Autonomous Nested Search for Hydrothermal Venting

Andrew Branch[1], Guangyu Xu[2], Michael V. Jakuba[2], Christopher R. German[2], Steve Chien[1], James C. Kinsey[2], Andrew D. Bowen[2], Kevin P. Hand[1], Jeffrey S. Seewald[2]

[1]Jet Propulsion Laboratory, California Institute of Technology
[2]Woods Hole Oceanographic Institution
Correspondence Author: andrew.branch@jpl.nasa.gov

Abstract

Ocean Worlds in the outer solar system represent one of the best chances for the discovery of extra-terrestrial life. Bodies, such as Europa and Enceladus, are thought to harbor liquid oceans, often encased in a thick icy shell. In order to further investigate these oceans, a new mission concept needs to be developed, a submersible craft. This vehicle would be required to traverse the icy shell and travel hundreds or even thousands of kilometers to survey the ocean below. In doing this, the vehicle might be out of contact for weeks or months at a time. The vehicle must be able to autonomously detect, locate, and study features of interest. One potential target is hydrothermal venting, due to their unique ecosystems on Earth. We have developed an autonomous, nested search strategy to locate sources of hydrothermal venting based on currently used methods. To test this search technique a simulation environment was developed using a hydrothermal plume dispersion simulation and a vehicle model. We show the effectiveness of the search method in this environment.

Introduction

At least eight bodies in our solar system are thought to harbor liquid oceans. In some cases, such as Europa and Enceladus, this ocean is perhaps habitable and encased in an icy shell kilometers thick [National Aeronautics and Space Administration 2018]. In order to explore these worlds new mission concepts must be developed using penetrating, submersible vehicles. A notional mission concept for such a submersible, outlined in Figure 1, contains four main components, an orbiting communications relay, a surface antenna, an under-ice base station, and a submersible vehicle.

In order to facilitate ice shell transit, the vehicle needs to be small (particularly in cross sectional area). The long mission duration (could take well over a year to melt through the icy shell and a one year mission) requires a low power vehicle, limiting the types of instruments on board. While the vehicle would ideally travel hundreds to thousands of kilometers distant from the base station, the submersible would need to return close to the base station to transfer data – with data subsequently relayed from the base station, through the surface antenna to the orbiter for eventual return to Earth. The radiation environment near the target body could preclude the use of an orbiting communication relay, instead relying on a relay in an eccentric Jovian orbit, in the case of Europa, increasing the time between communication windows from daily to the order of months. When the submersible is away from the base station it would be unable to communicate with Earth. Therefore, while making journeys further and further away from the base station, the submersible might be operating days or weeks without contact. During this time the submersible would be required to autonomously detect, locate, and study a specific feature of interest.

Hydrothermal venting is one potential target for a submersible mission. Evidence for hydrothermal activity has been found on one Ocean World, Enceladus [Waite et al. 2017]. On Earth, these geological phenomena harbor unique ecosystems and are potentially critical to the origin of life. Similar vents on Ocean Worlds could be the best chance at extra-terrestrial life in our Solar System. Hydrothermal vent localization on Earth is often done with a non-autonomous three-phase nested search [German et al. 2008]. [Yoerger et al. 2007a] demonstrates this strategy in a number of cruises. [Yoerger et al. 2007b] presents a method to autonomously revisit areas of interest after the primary mission is completed, however this requires humans to develop the primary mission. This method was used in the field multiple times. [Jakuba and Yoerger 2008] uses occupancy grid mapping in order to localize vents. [Saigol et al. 2010] uses a belief-maximization algorithm to find an object of interest in simulation. [Ferri, Jakuba, and Yoerger 2010] uses a trigger based approach in order to gather higher resolution data in areas of strong sensor readings. [Tian et al. 2014] uses an autonomous planning algorithm based on the chemical plume tracing of moths instead of a nested search strategy. Each of these approaches were tested in idealized simulation environments or with deployment data, which does not allow for testing of fully autonomous planning algorithms.

We have developed a fully autonomous nested search strategy for the localization of hydrothermal vents based on the three-phase nested search commonly used in the field. In order to test this approach we have developed a simulation environment using FVCOM, an existing ocean circulation model and a vehicle model.

The rest of the paper is organized as follows. First we discuss the structure of hydrothermal venting. Then we discuss the simulation environment used to test our approach. We outline the approach itself and the experimental setup.

Finally we discuss the results and what future work.

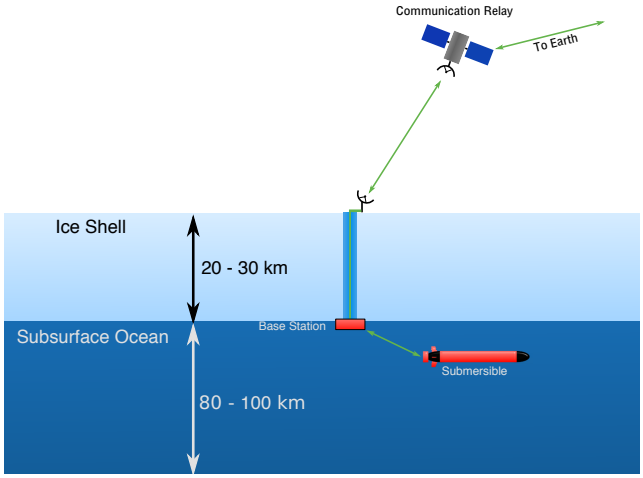

Figure 1: Notional Europa submersible mission showing the communication pathway from the submersible vehicle to Earth. Approximate ice thickness and ocean depth are labeled.

Hydrothermal Venting

Hydrothermal venting produces a plume which can be traced back to the source. The structure of the plume is shown in Figure 2. Hydrothermal fluid exiting the vent is less dense than the surrounding water, resulting in the formation of a buoyant plume. Due to entrainment this plume is continuously dilution the by ambient water column and expands from ~10 cm at the vent source to ~100 m once it reaches equilibrium. Upon reaching equilibrium, the plume expands horizontally — 10 to 100s of kilometers — to form the non-buoyant plume [German and Von Damm 2006]. The non-buoyant plume height is a function of the properties of the hydrothermal vent fluid as well as the surrounding water column [Turner 1979]. In the Pacific this height is approximately observed at 150 m above the seafloor, while in the Atlantic the plume height is closer to 300 m [Speer and Rona 1989].

Hydrothermal plumes are the main source of information when localizing venting. However, turbulent flow, differing vent types and strengths, and an unknown number of sources increase the difficulty of determining the plume source. Three primary sensors are used in the detection of hydrothermal plumes: temperature, optical backscatter [Baker, German, and Elderfield 1995; Baker and German 2004], and a chemical sensor such as oxidation-reduction potential [Nakamura et al. 2000]. These sensors are also good candidates for inclusion on a submersible mission to an Ocean World due to their compact form factor and low power consumption.

Simulation

A simulation environment was developed, using a hydrothermal plume dispersion simulation and a vehicle

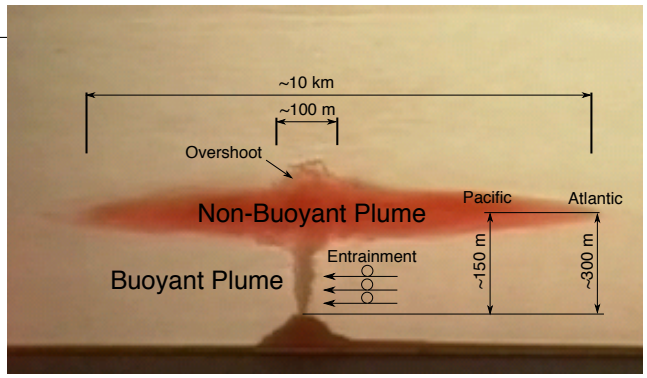

Figure 2: Model of a hydrothermal plume performed in an aquarium tank. The buoyant and non-buoyant components of the hydrothermal vent plume are labeled with approximate scales. Image courtesy of C. German, SOC UK

model. A numerical simulation of hydrothermal plume dispersion is performed using FVCOM, an ocean-circulation model, at Axial Seamount on the Juan de Fuca Ridge. The abundant lava supply to Axial supports vigorous hydrothermal systems and frequent volcanic activities, which have drawn extensive on-going scientific research that makes Axial one of the best-studied seamount on this planet.

FVCOM is a finite-volume, time and density-dependent, three-dimensional, ocean circulation model [Chen, Liu, and Beardsley 2003]. The unstructured grid employed in FVCOM supports grid size variation, therefore, proves efficient for the simulation of motion over a broad range of length scales. In addition, FVCOM supports the use of large-scale ocean circulation and tidal model outputs as open boundary forcing to drive flow across a broad range of frequencies inside the model domain [Zheng and Weisberg 2012].

Our model domain covers 300 by 300 km, centered on the Axial Seamount caldera and is open to flow across all four sides of that region. Horizontal resolution varies from 200 m within a 10 by 10km region enclosing Axial's caldera to 10km at the domain's boundary. The vertical dimension utilizes a uniform sigma-coordinate system with 127 layers, covering the full water column. This results in a ~12 m layer thickness above Axial's summit. The duration of the simulation is 58 days with 1 hour resolution. The 3-hourly sampled, 1/12.5° horizontal resolution, global reanalysis outputs of the HYbrid Coordinate Ocean Model (HYCOM) are used to construct the initial stratification profiles and open boundary forcing. Because HYCOM does not include ocean tides, we superimpose the tidal elevation and velocity predicted by the OSU Tidal Inversion onto the HYCOM outputs when constructing the open boundary forcing. We also add surface wind forcing and heat flux from 1-hourly sampled National Centers for Environmental Prediction (NCEP) Climate Forecast System Reanalysis (CFSR) outputs. We apply a linear ramp to bring open boundary and surface forcing from zero to full value over an initial four simulation days. Lastly, we add a seafloor heat source of 1 GW at the center (0,0) of the model domain inside Axial's caldera, which is turned on

after the initial four simulation days. The model output consists of current, temperature, salinity, and a passive tracer, dye, which is released at the vent source. This tracer has a value range of $[0, 1000]$. After 30 days the tracer content in a 20 by 20 km region surrounding the vent source reaches a quasi-steady state. In a 50 by 50 km region surrounding the vent source no quasi-steady state is reached before the end of the simulation.

The simulated vehicle has four degrees-of-freedom, surge, sway, heave, and yaw. A proportional controller allows the vehicle to navigate to a specified location. Simulated sensors are used to measure temperature, salinity, the passive tracer, vehicle depth, and distance to seafloor at a fixed interval. The position of the vehicle is assumed to be known at all times. Currently a chemical sensor, such as oxidation-reduction potential, and vehicle resources, such as energy and data capacity, are not modeled.

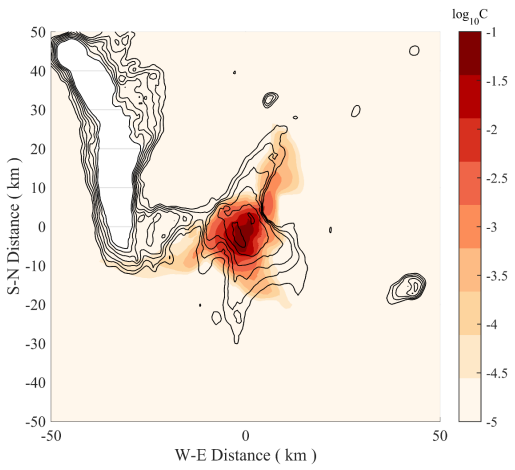

Figure 3: Snapshot taken at 1400 m depth on Mar 1, 2011 00:00 UTC of the simulated concentration (normalized by the source value) of a neutrally buoyant tracer originating from a hydrothermal vent source of 1 GW heat flux located inside the caldera of Axial Seamount at coordinate center. The global-simulation results of HYCOM and OSU Tidal Inversion for the period of Feb-Mar 2011 were used to drive flow inside the domain from its four boundaries.

Spatial Nested Search Strategy

The goal of this algorithm is, given a vehicle's starting location, produce a control strategy that results in locating the vent source. The vent source is considered found when the area around the vent has been surveyed at a specified resolution. In this case a resolution of 200 m was selected to match the resolution of the hydrothermal plume dispersion model at the vent source.

Before we can search for hydrothermal venting, we must have some method for detecting plumes. Ideally this would involve modeled sensors for temperature, optical backscatter, and oxidation reduction potential. However, currently we only use the passive tracer in the model as a direct measure of hydrothermal plume. This is an area of future improvement.

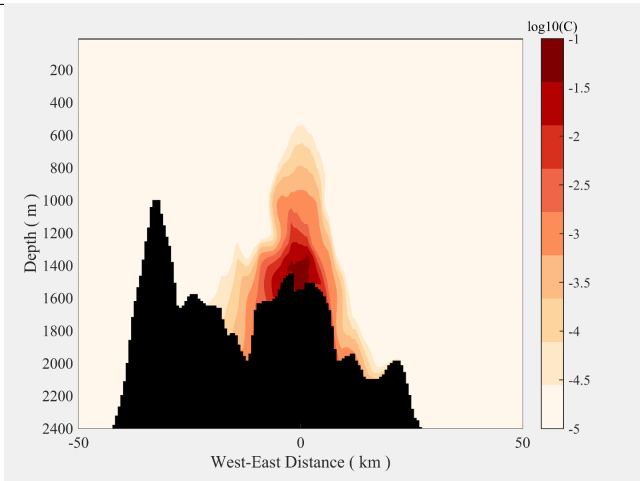

Figure 4: Snapshot taken along a W-E transect across the center of the model domain on Mar 1, 2011 00:00 UTC of the simulated concentration (normalized by the source value) of a neutrally buoyant tracer originating from a hydrothermal vent source of 1000 MW heat flux located inside the caldera of Axial Seamount at coordinate center. The global-simulation results of HYCOM and OSU Tidal Inversion for the period of Feb-Mar 2011 were used to drive flow inside the domain from its four boundaries.

The search algorithm operates as follows and is outlined in Algorithm 1. A spiral is initiated at the start location. The horizontal spacing of this spiral is manually selected to be the expected size of the feature in question. This insures features of the expected size are seen during this initial survey. During this spiral the vehicle completes vertical profiles through the extent of the water column. When the max plume detection value of a single profile exceeds the specified threshold, $plume\_thresh$ in Algorithm 1, the second phase of surveys begins. The height of the detected feature, $plume\_height$, is determined by binning the data from the vertical profile, $profile\_data$, at a 10 m resolution and selecting the bin with the largest average value. The subsequent surveys are performed at a depth of $plume\_height$.

During the second phase of surveys, the search space is partitioned into bins, $survey\_bins$ equal in size to the spacing of the initial lawnmower survey, $initial\_spacing$. These bins are separated into four quadrants centered on the corner of the bin closest to the location of the plume detection. A dynamic lawnmower pattern is executed in each of the four quadrants. The spacing of the lawnmower pattern, $initial\_spacing$, is once again specified beforehand. Each track line of the lawnmower pattern consists of sections with length equal to the spacing. At least $min\_sections$ sections will be completed for each track line. If $sections\_limit$ sections have average plume strength below some specified threshold and the sections have monotonically decreasing average plume strengths then the track line is completed and the next track line is commenced. $min\_sections$ and

$sections\_limit$ are search parameters specified beforehand. If the maximum value of the entire track line is less than $plume\_thresh$ then the lawnmower survey is completed and the next begins. The data from each dynamic lawnmower is binned into $survey\_bins$.

An example dynamic lawnmower is shown in Figure 5 and outlined in Algorithm 1. The plot is subdivided into track line sections. The average plume strength is listed in each section; a green background indicates that the average plume strength is greater than the specified threshold, $plume\_thresh$. Two boundaries to the survey are shown. Upon reaching the right-most boundary, the vehicle completes the current trackline. The boundaries correspond to the shared edges of the four quadrants during the search process.

Upon the completion of each dynamic lawnmower, local maxima of $survey\_bins$ are found and a nested lawnmower survey begin. An example of this process is shown in Figure 6. The local maximum, shown in green, and its neighbors are subdivided into smaller bins with one-third the side length of their parents. A lawnmower with spacing equal to one-third that of the previous lawnmower survey and with track lines centered on each row of nested bins is initiated. The new nested lawnmower survey covers the local maximum and all surrounding neighbors. If multiple local maxima have been found, they are prioritized on plume strength. This process repeats recursively until survey spacing of $final\_spacing$, is reached. If no maxima are found, or all local maxima have been exhausted, then the dynamic lawnmowers resume. After all dynamic lawnmower surveys are completed the spiral is resumed. Another set of dynamic lawnmowers is started if a plume is detected outside of the previously searched area.

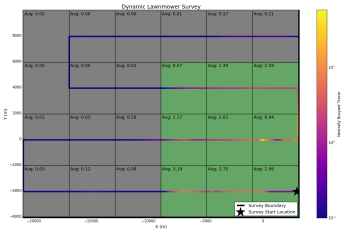

Figure 5: Plot showing an example dynamic lawnmower survey. The survey area is partitioned into regions representing sections of each track line. Regions shaded green have an average plume strength over the specified threshold. The average value is labeled in the upper left corner of each region. The two survey boundaries are thick black lines and the starting location is marked with a black star.

Experiment

121 scenarios were completed with the vehicle starting location uniformly varied between $x = [-30000, 30000]$ and $y = [-30000, 30000]$ at intervals of 6000m. Due to the nature of the algorithm and the location of the vent at $(0, 0)$ it is likely that the vehicle will pass directly over the vent source

Algorithm 1 Autonomous Nested Search

procedure NESTED_SEARCH
- $plans \leftarrow$ empty stack
- $visited \leftarrow$ empty set
- $plans.push(spiral)$
- $survey\_bins \leftarrow$ bins of size $inital\_spacing$
- **while** $plans.size > 0$ and not timed out **do**
  - Execute or Continue $plans.top()$
  - **if** executing spiral **then**
    - Wait until end of vertical profile
    - $profile\_data \leftarrow$ Get data from profile
    - $d \leftarrow max(profile\_data)$
    - **if** $d >= plume\_thresh$ and $d.location$ not explored **then**
      - $bins \leftarrow profile\_data$ binned at 10 meters and averaged
      - $plume\_height \leftarrow max(bins).height$
      - $(x, y) \leftarrow$ bin corner closest to $d.position$
      - $plans.push(dynamic\_lawnmower(x, y, plume\_height, 0, \pi/2))$
      - $plans.push(dynamic\_lawnmower(x, y, plume\_height, -\pi, \pi/2))$
      - $plans.push(dynamic\_lawnmower(x, y, plume\_height, -\pi, -\pi/2))$
      - $plans.push(dynamic\_lawnmower(x, y, plume\_height, 0, -\pi/2))$
      - Execute $plans.top()$
  - **else**
    - **while** $plans.top()$ is not completed **do**
      - Wait
    - $maxima \leftarrow get\_bin\_maxima(survey\_bins)$
    - sort $maxima$
    - **for** bin in $maxima$ **do**
      - **if** bin not in $visited$ **then**
        - Partition bin and $bin.neighbors()$
        - $visited.add(bin)$
        - $plans.push(nested\_lawnmower(bin))$
  - **while** $plans.size > 0$ and $plans.top()$ is complete **do**
    - $f \leftarrow plans.pop()$
    - **if** $f.spacing < final\_spacing$ and f contains vent source **then**
      - **return** Success
- **return** Failure

procedure EXECUTE_DYNAMIC_LAWNMOWER(x, y, h , along_track, across_track)
- $start\_x \leftarrow x + cos(along\_track) * inital\_spacing/2$
- $start\_y \leftarrow y + sin(across\_track) * inital\_spacing/2$
- Go to $(start\_x, start\_y, h)$
- $curr\_track \leftarrow 0$
- $curr\_section \leftarrow 0$
- $completed \leftarrow False$
- $section\_data \leftarrow$ empty list
- **while** not $completed$ **do**
  - Do next section on current track
  - $section\_data[curr\_section] \leftarrow$ Get data from last section
  - $curr\_section \leftarrow curr\_section + 1$
  - **if** $curr\_section >= min\_sections$ **then**
    - **if** $avg(section\_data[i]) < plume\_thresh$ for last $sections\_limit$ sections and monotonically decreasing **then**
      - $curr\_track \leftarrow curr\_track + 1$
      - **if** $max(section\_data) < plume\_thresh$ **then**
        - $completed \leftarrow True$
        - $section\_data \leftarrow$ empty list

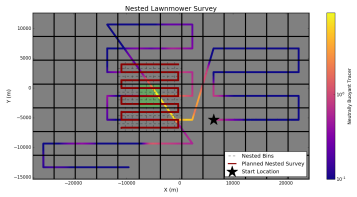

Figure 6: Plot showing an example of the planning process for nested lawnmowers. The search space is divided into square bins with sides equal to the lawnmower spacing. Upon finding a local maximum bin, the bin and all its neighbors are subdivided into nested bins of one-third the side length. A lawnmower pattern is then executed such that each track line is centered on a line of bins. The vehicle path and observed tracer is plotted. The planned nested lawnmower is show in dark red. The starting location is marked with a black star.

if the start location x and y are multiples of 1000. To mitigate this, a uniformly random value between $[-1500, 1500]$ was added to the x and y values of the starting location. The simulated vehicle has a horizontal and vertical velocity of 1 m/s. The vehicle samples the model at 0.2 hz. The plume detection threshold was set to 0.5. The initial spiral spacing was set to 5000 m and the initial dynamic lawnmower spacing was set to 4000 m.

Results

87% of the simulation scenarios successfully found the vent location within 28 days. Figure 7 shows the time each run took to successfully find the vent in black. The runs that failed to find the vent are shown in red. Plot (a) shows the total time while plots (b), (c), and (d) show the time spent on the spiral survey, dynamic lawnmower surveys, and nested lawnmower surveys respectively. Figures 8, 9, and 10 show an example run plotting a top down view and a 3d view of the passive tracer (dye) value from the model, and a top down view of the survey types during the run, respectively.

We see slight correlation between the distance and total time on successful runs. When this is broken down into the different stages of the algorithm we see this correlation stronger within the spiral surveys while not at all in the lawnmower surveys. No correlation is seen between the failed surveys and the distance from the vent, indicating that the cause of the failure is not related to distance. This method does not have a set distance in which it is feasible, starting further from the vent location would only require longer search times. Search times can be minimized by selecting appropriate values for the survey spacing parameters.

Future Work

Currently, the vehicle simulation is rudimentary. Realistic models for sensors such as temperature, optical backscatter, and chemical sensors can be developed. Vehicle resources such as power and data capacity can be implemented. Finally, the vehicles motion model can be improved by evecting the vehicle according to the currents in the model.

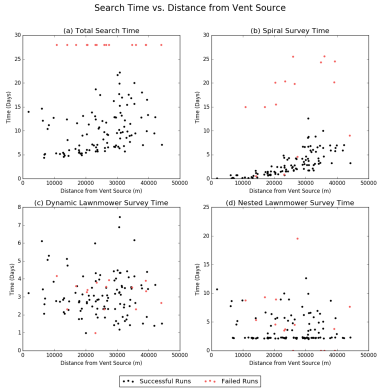

Figure 7: Plots showing the time taken to find the vent source compared to the distance from the vent source. The runs that successfully find the vent within 28 days are show in black. The failed runs are shown in red. Panel (a) plots the total time spent during the search. Panel (b), (c), and (d) decompose the time into the spiral survey, dynamic lawnmower surveys, and nested surveys respectively

The planning method has many areas which could use further investigation. Improved search in the vertical direction would insure that the vehicle maintains contact with the strongest part of the plume. The temporal variation of the plume is currently not accounted for by the search algorithm, potentially misleading the search. Other geometric search patterns and other search strategies such as gradient search or biologically inspired approaches can be implemented and tested. Vehicle resource considerations can be incorporated into the planner. More intelligent path planning can be implemented to reduce resource consumption while performing multiple surveys. Hydrothermal activity is one potential target for a submersible, investigation into other targets and the development of a search approach capable of prioritizing multiple target types would be beneficial.

The data volume collected by the vehicle far exceeds the communication throughput capabilities back to Earth. A method of summarizing the data collected for return to Earth needs to be developed. A number of spacecraft have implemented systems for this purpose. The Autonomous Sciencecraft Experiment used onboard science algorithms to summarize, delete, and prioritize data for downlink [Chien et al. 2005]. The onboard product generation for the Earth Observing-1 mission serves as a predecessor to the proposed HyspIRI Intelligent Payload Module [Chien et al. 2013]. The Mars Exploration Rover's (MER) WATCH system processes imagery to detect dust devils and send summarized data products to Earth [Castano et al. 2008]. The AEGIS system processes onboard imagery to autonomously retarget science instruments on the Mars Science Laboratory [Estlin et al. 2014] and MER [Estlin et al. 2012].

More simulation runs varying search parameters such as starting location, plume detection threshold, and survey spacing would allow for a better understanding of the presented search strategy. Another plume dispersal model, ei-

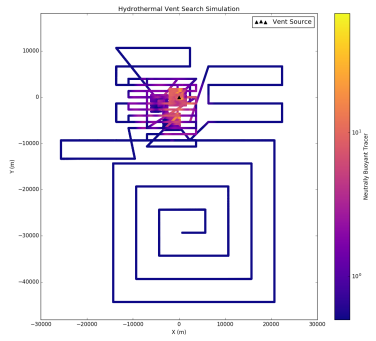

Figure 8: Top down plot showing the passive tracer (dye) as seen by the vehicle from a scenario starting at x=710, y=-29337. The vent source location is shown as a black triangle at (0,0).

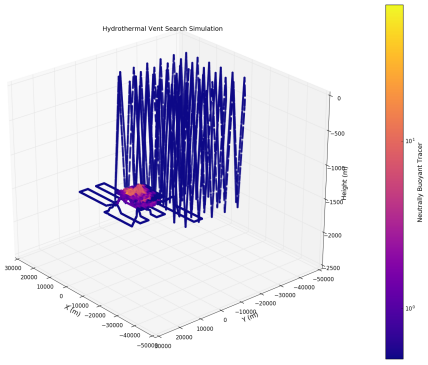

Figure 9: 3d view plot showing the passive tracer (dye) as seen by the vehicle from a scenario starting at x=710, y=-29337.

ther of a different region or with different plume parameters could be developed. Real world tests in well studied areas such as Axial Seamount would further validate the approach.

Conclusion

We have developed a realistic simulation environment in which to test search strategies for the localization of hydrothermal venting. We developed an autonomous nested search based on the current manual three-phase search method [German et al. 2008]. Search parameters, such as survey resolution and search location, allow for manual fine tuning of the search process based on the observed data, allowing for a human-in-the-loop model when possible. We performed 121 scenarios with varying start locations, of which 87% were able to successfully find the hydrothermal vent within 28 days.

Acknowledgments

Portions of this work were performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

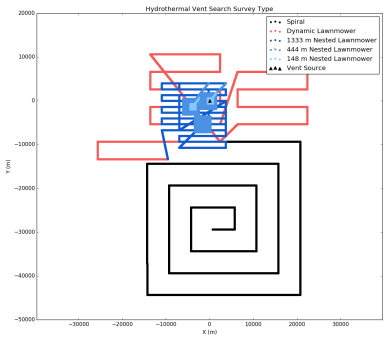

Figure 10: Plots showing the types of surveys performed on a scenario starting at x=710, y=-29337. The spiral survey is shown is black, the dynamic lawnmower surveys are red, and the nested lawnmower surveys are differing shades of blue with darker shades as surveys with larger spacing. The vent source location is shown as a black triangle at (0,0).

References

Baker, E. T., and German, C. R. 2004. On the global distribution of hydrothermal vent fields. *Mid-ocean ridges* 245–266.

Baker, E. T.; German, C. R.; and Elderfield, H. 1995. Hydrothermal plumes over spreading-center axes: Global distributions and geological inferences. *Seafloor hydrothermal systems: Physical, chemical, biological, and geological interactions* 47–71.

Castano, A.; Fukunaga, A.; Biesiadecki, J.; Neakrase, L.; Whelley, P.; Greeley, R.; Lemmon, M.; Castano, R.; and Chien, S. 2008. Automatic detection of dust devils and clouds on mars. *Machine Vision and Applications* 19(5-6):467–482.

Chen, C.; Liu, H.; and Beardsley, R. C. 2003. An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Application to coastal ocean and estuaries. *Journal of Atmospheric and Oceanic Technology* 20(1):159–186.

Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davis, A.; Mandl, D.; Trout, B.; Shulman, S.; et al. 2005. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication* 2(4):196–216.

Chien, S.; Mclaren, D.; Tran, D.; Davies, A. G.; Doubleday, J.; and Mandl, D. 2013. Onboard product generation on earth observing one: a pathfinder for the proposed hyspiri mission intelligent payload module. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 6(2):257–264.

Estlin, T. A.; Bornstein, B. J.; Gaines, D. M.; Anderson,

R. C.; Thompson, D. R.; Burl, M.; Castano, R.; and Judd, M. 2012. Aegis automated science targeting for the mer opportunity rover. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3(3):50.

Estlin, T.; Gaines, D.; Bornstein, B.; Schaffer, S.; Tompkins, V.; Thompson, D. R.; Altinok, A.; Anderson, R. C.; Burl, M.; Castaño, R.; et al. 2014. Automated targeting for the msl rover chemcam spectrometer. In *12th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, 17–19.

Ferri, G.; Jakuba, M. V.; and Yoerger, D. R. 2010. A novel trigger-based method for hydrothermal vents prospecting using an autonomous underwater robot. *Autonomous Robots* 29(1):67–83.

German, C., and Von Damm, K. 2006. Hydrothermal processes. *Treatise on geochemistry* 6:181–222.

German, C. R.; Yoerger, D. R.; Jakuba, M.; Shank, T. M.; Langmuir, C. H.; and Nakamura, K.-i. 2008. Hydrothermal exploration with the autonomous benthic explorer. *Deep Sea Research Part I: Oceanographic Research Papers* 55(2):203–219.

Jakuba, M., and Yoerger, D. R. 2008. Autonomous search for hydrothermal vent fields with occupancy grid maps. In *Proc. of ACRA*, volume 8, 2008.

Nakamura, K.; Veirs, S.; Sarason, C. P.; McDuff, R. E.; Stahr, F.; Yoerger, D. R.; and Bradley, A. M. 2000. Electrochemical signals in rising buoyant plumes and tidally oscillating plumes at the main endeavour vent field, juan de fuca ridge. *EOS, Transactions of the American Geophysical Union* 81(48).

National Aeronautics and Space Administration. 2018. Ocean worlds.

Saigol, Z.; Dearden, R.; Wyatt, J.; and Murton, B. 2010. Belief change maximisation for hydrothermal vent hunting using occupancy grids. In *Proceedings of the Eleventh Conference Towards Autonomous Robotic Systems (TAROS-10)*, 247–254.

Speer, K. G., and Rona, P. A. 1989. A model of an atlantic and pacific hydrothermal plume. *Journal of Geophysical Research: Oceans* 94(C5):6213–6220.

Tian, Y.; Zhang, A.; Li, W.; Yu, J.; Li, Y.; and Zeng, J. 2014. A behavior-based planning strategy for deep-sea hydrothermal plume tracing with autonomous underwater vehicles. In *OCEANS 2014-TAIPEI*, 1–10. IEEE.

Turner, J. S. 1979. *Buoyancy effects in fluids*. Cambridge University Press.

Waite, J. H.; Glein, C. R.; Perryman, R. S.; Teolis, B. D.; Magee, B. A.; Miller, G.; Grimes, J.; Perry, M. E.; Miller, K. E.; Bouquet, A.; Lunine, J. I.; Brockwell, T.; and Bolton, S. J. 2017. Cassini finds molecular hydrogen in the enceladus plume: Evidence for hydrothermal processes. *Science* 356(6334):155–159.

Yoerger, D. R.; Bradley, A. M.; Jakuba, M. V.; Tivey, M. A.; German, C. R.; Shank, T. M.; and Embley, R. W. 2007a. Mid-ocean ridge exploration with an autonomous underwater vehicle.

Yoerger, D. R.; Jakuba, M.; Bradley, A. M.; and Bingham, B. 2007b. Techniques for deep sea near bottom survey using an autonomous underwater vehicle. *The International Journal of Robotics Research* 26(1):41–54.

Zheng, L., and Weisberg, R. H. 2012. Modeling the west florida coastal ocean by downscaling from the deep ocean, across the continental shelf and into the estuaries. *Ocean Modelling* 48:10 – 29.